Solving differential games by methods for finite-dimensional saddle-point problems

Pavel Dvurechensky, Vladimir Spokoiny
Weierstrass Institute for Applied Analysis and Stochastics
{pavel.dvurechensky,spokoiny}@wias-berlin.de

Yurii Nesterov
COREUniversite catholique de Louvain
yurii.nesterov@uclouvain.be

Abstract

Several recent papers propose a perspective on Deep Neural Networks through the lens of optimal control theory. This leads to a new analysis, for example of the stability of backpropagation. On the other hand, saddle-point problems and variational inequalities gained a new interest in the deep learning community, e.g. in application to Generative Adversary Networks. The field of differential games combines saddle-point problems and optimal control in order to study optimal control problems in the presence of conflict or uncertainty. In this work we consider convex-concave differential games and propose an algorithm to find the optimal strategies. The algorithm is based on the construction of dual problem, which is finite-dimensional problem, and applying a primal-dual gradient method. A careful analysis is used to prove the rates in terms of the solution to the original problem.[1]

1 Introduction

Several recent papers propose a perspective on Deep Neural Networks through the lens of optimal control theory [4, 5]. This leads to a new analysis, for example of the stability of backpropagation. On the other hand, saddle-point problems and variational inequalities gained a new interest in the deep learning community, e.g. in application to Generative Adversary Networks [3, 6]. The field of differential games combines saddle-point problems and optimal control in order to study optimal control problems in the presence of conflict or uncertainty. In this work we consider convex-concave differential games and propose an algorithm to find the optimal strategies.

In the last years we can observe an increasing interest to the primal-dual subgradient methods. This line of research, started in [8], leads to the special methods, which allow to reconstruct approximate solution to a *conjugate* problem. In order to do this, methods need to get an access to the internal variables of the oracle. Therefore, all these methods are problem-specific.

This approach works well when the primal and conjugate problems have different level of complexity. For example, we can have a primal minimization problem of very high dimension, with very simple objective function and basic feasible set, and a small number of linear equations. Introducing Lagrange multipliers for these linear constraints, we can pass to the conjugate (dual) problem[2], which

[1]The full work was published as [2].

[2]Since the objective and the feasible set of our problem are simple, very often this can be done in an explicit form.

Abstract for: Smooth Games Optimization and Machine Learning Workshop (NIPS 2018), Montréal, Canada.

has good chances to be simple in view of its small dimensional. The only delicate problem is the reconstruction of the primal variables from the minimization process, which we run in the conjugate space.

In [1] this approach was applied to the problems of Optimal Control with convex constraints for the end point of the trajectory. These constraints were treated by linear operators from infinite-dimensional space of variables (control) to a finite-dimensiomal space of phase variables. It was shown, that an appropriate optimization process in the latter space can generate also nearly optimal sequence of controls (functions of time). Moreover, this technique was supported by the worst-case complexity analysis.

In this work we do the next step in this direction. We consider an infinite-dimensional saddle-point problem, which variables (controls) must satisfy some linear equality constraints. We show that these constraints can be dualized by *finite-dimensional* multipliers. Moreover, it appears that the dual counterpart of our problem is again a saddle-point problem, but in a finite dimension (we call this problem *conjugate*). We show how to reconstruct the infinite-dimensional primal strategies from a special finite-dimensional scheme, which solves the conjugate problem.

2 Main results

Consider two moving objects with dynamics given by the following equations:

$$\dot{x}(t) = A_x(t)x(t) + B(t)u(t), \quad \dot{y}(t) = A_y(t)y(t) + C(t)v(t), \quad (x(0), y(0)) = (x_0, y_0). \tag{1}$$

Here $x(t) \in \mathbb{R}^n$, $y(t) \in \mathbb{R}^m$ are the phase vectors of these objects, $u(t)$ is the control of the first object (pursuer), and $v(t)$ is the control of the second object (evader). Matrices $A_x(t), A_y(t), B(t)$, and $C(t)$ are continuous and have appropriate sizes. The system is considered on the time interval $[0, \theta]$. Controls are restricted in the following way $u(t) \in P \subset \mathbb{R}^p$, $v(t) \in Q \subset \mathbb{R}^q \quad \forall t \in [0, \theta]$. We assume that P, Q are closed convex sets.

The goal of pursuer is to minimize the value of the functional:

$$F(u, v) + \Phi(x(\theta), y(\theta)) \stackrel{\text{def}}{=} \int_0^\theta \tilde{F}(\tau, u(\tau), v(\tau))d\tau + \Phi(x(\theta), y(\theta)). \tag{2}$$

The goal of the evader is the opposite. We need to find an optimal guaranteed result for each object, which leads to the problem of finding the saddle-point of the above functional. We assume the following:

- $u(\cdot) \in L_2([0, \theta], \mathbb{R}^p)$, and $v(\cdot) \in L_2([0, \theta], \mathbb{R}^q)$,
- the saddle-point in this class of strategies exists,
- function $F(u, v)$ is upper semi-continuous in v and lower semi-continuous in u,
- $\Phi(x, y)$ is continuous.

If we solve the first differential equation in (1), then we can express $x(\theta)$ as a result of application of the linear operator $\mathcal{B} : L_2([0, \theta], \mathbb{R}^p) \to \mathbb{R}^n$:

$$x(\theta) \stackrel{\text{def}}{=} \tilde{x}_0 + \mathcal{B}u, \tag{3}$$

and the linear operator $\mathcal{C}$ is defined in the same manner.

So our main problem of interest has the following form:

$$\min_{u \in P} \left[\max_{v \in Q} \{ F(u, v) + \Phi(x, y) : y = \tilde{y}_0 + \mathcal{C}v \} : x = \tilde{x}_0 + \mathcal{B}u \right] \tag{4}$$

In order to obtain a smooth saddle-point problem, let us introduce some assumptions.

A1 function $F(\cdot, v)$ is strongly convex for any fixed v with constant σ_{F_u} which does not depend on v, and function $F(u, \cdot)$ is strongly concave for any fixed u with constant σ_{F_v} which does not depend on u. Also assume that:

$$\|\nabla_u F(u, v_1) - \nabla_u F(u, v_2)\| \leq L_{uv} \|v_1 - v_2\|, \|\nabla_v F(u_1, v) - \nabla_v F(u_2, v)\| \leq L_{vu} \|u_1 - u_2\| \tag{5}$$

A2 $\Phi(\cdot, y)$ is strongly convex for any fixed y with constant $\sigma_{\Phi x}$ which doesn't depend on y and $\Phi(x, \cdot)$ is strongly concave for any fixed x with constant $\sigma_{\Phi y}$ which doesn't depend on x. Also assume that:

$$\|\nabla_x \Phi(x, y_1) - \nabla_x \Phi(x, y_2)\| \le L_{xy} \|y_1 - y_2\|, \|\nabla_y \Phi(x_1, y) - \nabla_y \Phi(x_2, y)\| \le L_{yx} \|x_1 - x_2\|, \tag{6}$$

and

$$\|\nabla_x \Phi(x_1, y) - \nabla_x \Phi(x_2, y)\| \le L_{xx} \|x_1 - x_2\|, \|\nabla_y \Phi(x, y_1) - \nabla_y \Phi(x, y_2)\| \le L_{yy} \|y_1 - y_2\|. \tag{7}$$

Note that assumptions **A1**, **A2** imply that the level sets of functions $F(u, v), \Phi(x, y)$ are closed convex and bounded. We construct the conjugate problem for (4), which is

$$\begin{aligned} \min_\lambda \max_\mu \{ \quad & \min_{u \in P} \max_{v \in Q} [F(u, v) - \langle \mu, \mathcal{B}u \rangle + \langle \lambda, \mathcal{C}v \rangle] \\ & + \min_x \max_y [\Phi(x, y) + \langle \mu, x \rangle - \langle \lambda, y \rangle] - \langle \mu, \tilde{x}_0 \rangle + \langle \lambda, \tilde{y}_0 \rangle \ \}. \end{aligned} \tag{8}$$

Here $\lambda \in \mathbb{R}^n$ and $\mu \in \mathbb{R}^m$.

We assume that problems

$$\psi_1(\lambda, \mu) = \min_{u \in P} \max_{v \in Q} [F(u, v) - \langle \mu, \mathcal{B}u \rangle + \langle \lambda, \mathcal{C}v \rangle], \tag{9}$$

$$\psi_2(\lambda, \mu) = \min_{x} \max_{y} [\Phi(x, y) + \langle \mu, x \rangle - \langle \lambda, y \rangle] \tag{10}$$

are simple, which means that they can be solved efficiently or in a closed form. Note that the conjugate problem is finite-dimensional. Using assumptions **A1**, **A2**, the fact that closed convex bounded set in Hilbert space is compact in weak topology, the fact that $F(u, v)$ is upper semi-continuous in v and lower semi-continuous in u, we conclude that the saddle points in the problems (9), (10) do exist for all $\lambda \in \mathbb{R}^n$ and $\mu \in \mathbb{R}^m$.

Lemma 1. *Let Assumptions* **A1**, **A2** *hold. Then $\psi_1(\lambda, \mu)$ and $\psi_2(\lambda, \mu)$ are smooth with partial gradients satisfying the Lipschitz condition.*

Next, we describe the setup for the algorithm. Let us introduce prox-function $d_\lambda(\lambda) = \frac{\sigma_\lambda}{2} \|\lambda\|^2$, where we use Euclidian norm. Function $d_\lambda(\lambda)$ is strongly convex in this norm with convexity parameter σ_λ. For the variable μ we introduce prox-function $d_\mu(\mu) = \frac{\sigma_\mu}{2} \|\mu\|^2$, which is strongly convex with convexity parameter σ_μ with respect to Euclidian norm. These prox-functions are differentiable everywhere. For any $\lambda_1, \lambda_2 \in \mathbb{R}^n$ we define Bregman divergence $\omega_\lambda(\lambda_1, \lambda_2) = \frac{\sigma_\lambda}{2} \|\lambda_1 - \lambda_2\|^2$ Let us choose $\bar{\lambda} = 0$ as the center of the space $\mathbb{R}^n$. Then we have $\omega_\lambda(\bar{\lambda}, \lambda) = d_\lambda(\lambda)$. For μ we introduce the similar settings.

Finding the saddle point (λ^*, μ^*) for conjugate problem (8) is equivalent to solving variational inequality

$$\langle g(\lambda, \mu), (\lambda - \lambda^*, \mu - \mu^*) \rangle \ge 0, \quad \forall \lambda, \mu, \tag{11}$$

where

$$g(\lambda, \mu) = (\nabla_\lambda \psi(\lambda, \mu), -\nabla_\mu \psi(\lambda, \mu)). \tag{12}$$

Let us choose some $\kappa \in (0, 1)$. Consider a space of $z \stackrel{\text{def}}{=} (\lambda, \mu)$ with the norm

$$\|z\| = \sqrt{\kappa \sigma_\lambda \|\lambda\|^2 + (1 - \kappa) \sigma_\mu \|\mu\|^2},$$

an oracle $g(z) = (\nabla_\lambda \psi(\lambda, \mu), -\nabla_\mu \psi(\lambda, \mu))$, a new prox-function $d(z) = \kappa d_\lambda(\lambda) + (1 - \kappa) d_\mu(\mu)$ which is strongly convex with constant $\sigma_0 = 1$. We define $W = \mathbb{R}^n \times \mathbb{R}^m$, Bregman divergence $\omega(z_1, z_2) = \kappa \omega_\lambda(\lambda_1, \lambda_2) + (1 - \kappa) \omega_\lambda(\mu_2, \mu_2)$ which has an explicit form of $\omega(z_1, z_2) = d(z_1 - z_2)$, and center $\bar{z} = (0, 0)$. Then we have $\omega(\bar{z}, z) = d(z)$. Note that the norm in the dual space is defined as

$$\|g\| = \sqrt{\frac{1}{\kappa \sigma_\lambda} \|g_\lambda\|^2 + \frac{1}{(1 - \kappa) \sigma_\mu} \|g_\mu\|^2}$$

Lemma 2. *Let Assumptions* **A1**, **A2** *be true, and* $\kappa = \frac{\sigma_\mu}{\sigma_\mu+\sigma_\lambda}$. *Then operator* $g(z)$ *defined in* (12) *is Lipschitz continuous:*

$$\|g(z_1) - g(z_2)\| \le L\,\|z_1 - z_2\|\,, \tag{13}$$

where L depends on the problem parameters.

We use the Dual Extrapolation method from [7] to solve (11)

1. Initialization: Fix $\beta = \frac{L}{\sigma_0}$. Set $s_{-1} = 0$.
2. Iteration ($k \ge 0$):
   Compute $x_k = T_\beta(\bar{z}, s_{k-1})$, (M1)
   Compute $z_k = T_\beta(x_k, -g(x_k))$,
   Set $s_k = s_{k-1} - g(z_k)$.

Here

$$T_\beta(z,s) = \arg\max_{x\in W}\{\langle s, x-z\rangle - \beta\omega(z,x)\}.$$

Similarly to [8], we can prove that method (M1) generates bounded sequence $\{z_i\}_{i\ge 0}$. Hence the sequences $\{\lambda_i\}_{i\ge0}, \{\mu_i\}_{i\ge0}$ are also bounded. Also since the saddle point in the problem (4) exists, there exists a saddle point (λ^*, μ^*) for conjugate problem (8). These arguments allow us to choose D_λ, D_μ such that $d_\lambda(\lambda_i) \le D_\lambda$, $d_\mu(\mu_i) \le D_\mu$ for all $i \ge 0$, which also ensure that (λ^*, μ^*) is an interior solution:

$$B_{r/\sqrt{\kappa\sigma_\lambda}}(\lambda^*) \subseteq W_\lambda \stackrel{\text{def}}{=} \{\lambda : d_\lambda(\lambda) \le D_\lambda\}\,, B_{r/\sqrt{(1-\kappa)\sigma_\mu}}(\mu^*) \subseteq W_\mu \stackrel{\text{def}}{=} \{\mu : d_\mu(\mu) \le D_\mu\}$$

for some $r > 0$. Then we have $z^* = (\lambda^*, \mu^*) \in \mathcal{F}_D \stackrel{\text{def}}{=} \{z \in W : d(z) \le D\}$ with $D = \kappa D_\lambda + (1-\kappa)D_\mu$ and $B_r(z^*) \subseteq \mathcal{F}_D$.

To characterize the quality of the solution, we define a function

$$\begin{aligned}\phi(u,x,v,y) = \min\nolimits_\lambda \max\nolimits_\mu\{ \quad & F(u,v) + \Phi(x,y) + \langle \mu, x - \tilde{x}_0 - \mathcal{B}u\rangle + \langle \lambda, \mathcal{C}v + \tilde{y}_0 - y\rangle : \\ & d_\lambda(\lambda) \le D_\lambda, d_\mu(\mu) \le D_\mu\}.\end{aligned} \tag{14}$$

Since $d_\lambda(\lambda^*) \le D_\lambda$, $d_\mu(\mu^*) \le D_\mu$, and the conjugate problem is equivalent to the initial one, we conclude that the initial problem is equivalent to the problem

$$\min_{u\in P, x\in X}\ \max_{v\in Q, y\in Y}\ \phi(u,x,v,y) \tag{15}$$

Let us introduce two auxiliary functions:

$$\xi(u,x) = \max_{v\in Q, y\in Y} \phi(u,x,v,y), \quad \eta(v,y) = \min_{u\in P, x\in X} \phi(u,x,v,y) \tag{16}$$

Note that $\xi(u,x)$ is convex, $\eta(v,y)$ is concave, and $\xi(u,x) \ge \phi(u^*,x^*,v^*,y^*) \ge \eta(v,y) \quad \forall u \in P, v \in Q, x \in X, y \in Y$, where $\phi(u^*,x^*,v^*,y^*)$ is the solution to (15).

Theorem 1. *Let Assumptions* **A1** *and* **A2** *be true, and* $\kappa = \frac{\sigma_\mu}{\sigma_\mu+\sigma_\lambda}$, L *be defined in Lemma 2. Let the points* $z_i = (\lambda_i, \mu_i), i \ge 0$ *be generated by method* (M1). *Let*

$$\hat{u}_{k+1} = \frac{1}{k+1}\sum_{i=0}^k u_i, \quad \hat{v}_{k+1} = \frac{1}{k+1}\sum_{i=0}^k v_i, \quad \hat{x}_{k+1} = \frac{1}{k+1}\sum_{i=0}^k x_i, \quad \hat{y}_{k+1} = \frac{1}{k+1}\sum_{i=0}^k y_i, \tag{17}$$

where (u_i, v_i), (x_i, y_i) *are the saddle points at the point* (λ_i, μ_i) *in* (9) *and* (10) *respectively. Then for functions* $\xi(u,x), \eta(v,y)$ *defined in* (16) *we have:*

$$\xi(\hat{u}_{k+1}, \hat{x}_{k+1}) - \eta(\hat{v}_{k+1}, \hat{y}_{k+1}) \le \frac{LD}{k+1} \tag{18}$$

Also the following is true:

$$\|\mathcal{B}\hat{u}_{k+1} + \tilde{x}_0 - \hat{x}_{k+1}\| \le \frac{LD\sqrt{\sigma_\mu}}{r(k+1)}, \quad \|\mathcal{C}\hat{v}_{k+1} + \tilde{y}_0 - \hat{y}_{k+1}\| \le \frac{LD\sqrt{\sigma_\lambda}}{r(k+1)}. \tag{19}$$

References

[1] O. Devolder, F. Glineur, and Y. Nesterov. Double smoothing technique for large-scale linearly constrained convex optimization. *SIAM Journal on Optimization*, 22(2):702–727, 2012.

[2] P. Dvurechensky, Y. Nesterov, and V. Spokoiny. Primal-dual methods for solving infinite-dimensional games. *Journal of Optimization Theory and Applications*, 166(1):23–51, Jul 2015.

[3] G. Gidel, H. Berard, P. Vincent, and S. Lacoste-Julien. A variational inequality perspective on generative adversarial networks. *arXiv:1802.10551*, 2018.

[4] E. Haber, L. Ruthotto, E. Holtham, and S.-H. Jun. Learning across scales - a multiscale method for convolution neural networks. In *AAAI Conference on Artificial Intelligence*, 2017. arXiv:1703.02009.

[5] J. Han and W. E. Deep learning approximation for stochastic control problems. *arXiv:1611.07422*, 2018.

[6] L. Mescheder, S. Nowozin, and A. Geiger. The numerics of gans. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 1825–1835. Curran Associates, Inc., 2017.

[7] Y. Nesterov. Dual extrapolation and its applications to solving variational inequalities and related problems. *Mathematical Programming*, 109(2-3):319–344, 2007. First appeared in 2003 as CORE discussion paper 2003/68.

[8] Y. Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical Programming*, 120(1):221–259, Aug 2009. First appeared in 2005 as CORE discussion paper 2005/67.